

**DYNAMIC MANAGEMENT OF COMPUTER  
WORKLOADS THROUGH SERVICE LEVEL OPTIMIZATION**

**ABSTRACT OF THE DISCLOSURE**

The inventive work load manager (WLM) dynamically responds to changes in workload characteristics. The WLM bases response on performance goals set by the administrator, and manipulates the resources to achieve these goals. The WLM receives performance information from performance monitors. The WLM examines the information from the performance monitors and compares the information with the goals using a Proportional Integral and Derivative controller. Based on the comparison, the WLM may increase, decrease, or leave unchanged, the resources devoted to an application. If the performance of an application is performance is lagging, e.g., if transactions are taking longer than the goal, then the WLM increases the entitlement. If an application is over-achieving, then the WLM will decrease its entitlement and allocate it to another application. The WLM uses an arbiter which allocates the resources between the different applications, based on the priorities of the applications. Consequently, the WLM creates a feedback loop between application performance and CPU entitlements.